GN Docket No. 14-28

BITAG

Broadband Internet Technical Advisory Group

Port Blocking

A BROADBAND INTERNET TECHNICAL ADVISORY GROUP
TECHNICAL WORKING GROUP REPORT

A Uniform Agreement Report

Issued:
August 2013

Executive Summary

The term "port blocking" refers to the practice of an Internet Service Provider (ISP) identifying Internet traffic by the combination of port number and transport protocol, and blocking it entirely. Port blocking thus affects the traffic associated with a particular combination of port number and transport protocol on that ISP, regardless of source or destination IP address. The practice can potentially prevent the use of particular applications altogether by blocking the ports those applications use. Port blocks can be deployed in a range of network locations, from where the ISP connects with other networks to datacenters and customer locations.

The Internet was built around the premise of an open and shared environment. Additionally, Internet standards assume all hosts on the global Internet can connect directly to each other, on any specified port number. The practical reality is that blocking of Internet port numbers, either in the short or long term, is a technique that has been used by both wireline and wireless network providers for various reasons for over a decade.

One of the original and enduring motivations for blocking ports is to prevent network attacks and abuse associated with particular application protocols. Some network and security administrators view port blocking as a critical tool for securing systems and information, and see it as part of the ISP's mission to manage the security risk to its users from theft and destruction of personal information, business records, and other critical electronic forms of information. TCP port 25, used for sending email, is an example of a port that is blocked by some operators to prevent network abuse – such as spam email.

Port blocking has also been used to enforce ISPs' terms of service. Likewise, port blocking was once viewed as a useful tool for managing capacity and bandwidth-intensive applications such as peer-to-peer file-sharing applications on enterprise and university networks. However, increased network capacity and a variety of developments in the application space have caused most residential ISPs to seek other ways of managing capacity. Finally, though rare, port blocking has at times been used to hinder competing applications, such as Voice over IP (VoIP).

Port blocking is among a set of tools and tactics (Network Address Translation (NAT) being the other major example) that can compromise the original intent of ports: to provide reliable local addresses so that end systems can manage multiple communications at once.

Port blocking can complicate application design and development and create uncertainty about whether applications will function properly when they are deployed. Port blocking can also cause applications to not function properly or "break" by preventing applications from using the ports they were designed to use. One of the outcomes of port blocking is an increase in the use of "port overloading." Port overloading is a tactic whereby application developers will design applications to use a common port, in order to minimize the chance of a port blocking practice impacting the usability of that application.

Importantly, it may not be obvious to Internet users why an application affected by port blocking is not working properly, because the application may simply be unable to connect or fail silently. If error messages are provided, they may not contain specific details as to the cause of the problem. Users may seek assistance from the ISP's customer service, online documentation, or other knowledgeable sources if they cannot diagnose the problem themselves. The fact that the problem could alternatively be caused by home networking equipment or a software-based port block complicates the process of diagnosis.

Users' ability to respond to port blocking depends on their technical sophistication and the extent to which workarounds are available. Overcoming port blocking may require the user to install a software update, change a configuration setting, request an opt-out from the ISP, or to upgrade their level of service (for example from residential to business). If these options are not available, or if users or customers lack the knowledge or willingness to pursue them, users may be prevented from using the blocked application altogether, or they may have to switch to a different application or a different network (from wireless to wireline, for example).

Because port blocking can affect how particular Internet applications function, its use has the potential to be anti-competitive, discriminatory, otherwise motivated by non-technical factors, or construed as such. As a result, the Broadband Internet Technical Advisory Group (BITAG) has a number of suggested practices when it comes to port blocking:

- **ISPs should avoid port blocking unless they have no reasonable alternatives available for preventing unwanted traffic and protecting users.** Further, if port blocking is deemed necessary, it should only be used for the purposes of protecting the implementing ISP's network and users. Port blocking should not be used for ongoing capacity management, to enforce non-security terms of service, or to disadvantage competing applications.
- **ISPs that can reasonably provide to their users opt-out provisions or exceptions to their port blocking policies should do so.** Whether opt-out provisions can be supported may depend on the particulars of the access network technology, the location port blocking is implemented in the network, administrative complexity, cost, and other factors.
- **ISPs should publicly disclose their port blocking policies.** The information should be readily available to both customers and non-customers alike, and should be as informative and concise as possible. For example, port blocking policies could be provided on the ISP's public facing website, on a page dedicated to summarizing or describing the respective ISP's network management practices.

  For persistent port blocks the information should include: (1) port numbers, (2) transport protocol (e.g., TCP or UDP), (3) the application(s) normally associated with the port(s), (4) the direction of the block – whether inbound or outbound, (5) a brief description of the reason(s) for the block, and (6) if opt-out provisions are available and how to request such.

- **ISPs should make communications channels available for feedback about port blocking policies.** Applications providers and consumers should have communications channels or other clear methods to discuss impacts caused by port blocking and to consider possible mitigations.
- **ISPs should revisit their port blocking policies on a regular basis and reassess whether the threats that required the port blocking rules continue to be relevant.** Some security threats are permanent and some are transitory or short-lived. Items such as spam prevention by blocking TCP port 25 from the customer are expected to last quite some time, while others such as blocks to prevent certain types of malicious software may be temporary.
- **Port blocking (or firewall) rules of consumers' devices should be user-configurable.** It is recommended that the documentation provided with each unit inform the consumer that port blocking or firewall rules have been implemented, which ports are blocked by default, and how consumers can modify those rules.

Table of Contents

1. About the BITAG

The Broadband Internet Technical Advisory Group (BITAG) is a non-profit, multi-stakeholder organization focused on bringing together engineers and technologists in a Technical Working Group (TWG) to develop consensus on broadband network management practices and other related technical issues that can affect users' Internet experience, including the impact to and from applications, content and devices that utilize the Internet.

The BITAG's mission includes: (a) educating policymakers on such technical issues; (b) addressing specific technical matters in an effort to minimize related policy disputes; and (c) serving as a sounding board for new ideas and network management practices. Specific TWG functions also may include: (i) identifying "best practices" by broadband providers and other entities; (ii) interpreting and applying "safe harbor" practices; (iii) otherwise providing technical guidance to industry and to the public; and/or (iv) issuing advisory opinions on the technical issues germane to the TWG's mission that may underlie disputes concerning broadband network management practices.

BITAG TWG reports focus primarily on technical issues. While the reports may touch on a broad range of questions associated with a particular network management practice, the reports are not intended to address or analyze in a comprehensive fashion the economic, legal, regulatory or public policy issues that the practice may raise.

The BITAG Technical Working Group and its individual Committees make decisions through a consensus process, with the corresponding levels of agreement represented on the cover of each report. Each TWG Representative works towards achieving consensus around recommendations their respective organizations support, although even at the highest level of agreement, BITAG consensus does not require that all TWG member organizations agree with each and every sentence of a document. The Chair of each TWG Committee determines if consensus has been reached. In the case there is disagreement within a Committee as to whether there is consensus, BITAG has a voting process with which various levels of agreement may be more formally achieved and indicated. For more information please see the BITAG Technical Working Group Manual, available on the BITAG website at www.bitag.org.

BITAG welcomes public comment. Please feel free to submit comments in writing via email at comments@bitag.org.

2. Issue Overview

The term "port blocking" refers to the practice of an Internet Service Provider (ISP) identifying Internet traffic by the combination of port number and transport protocol, and blocking it entirely. Port blocking thus affects the traffic associated with a particular combination of port number and transport protocol on an ISP, regardless of source or destination IP address. The practice can potentially prevent the use of particular applications altogether by blocking the ports those applications use. (Internet traffic may, of course, be treated in other ways—for example by redirecting it, rate limiting it, or changing its QoS classification—but such treatments are outside the scope of this report.)

Port blocking has been in use at various times by both wireline and wireless network operators for over a decade. One of the original and enduring motivations for blocking ports is to prevent network attacks and abuse associated with particular application protocols. Port blocking has also been used to enforce ISPs' terms of service – blocking inbound[1] port 80 for users[2] whose residential contracts prohibit them from running web servers, for example. This practice has become less common but is still in use by some operators. Likewise, port blocking was once viewed as a useful tool for managing capacity and bandwidth-intensive applications such as peer-to-peer file-sharing applications on enterprise and university networks. However, increased network capacity and a variety of developments in the application space have caused most residential ISPs to seek other ways of managing capacity. Finally, though rare, port blocking has at times been used by network operators to hinder competing applications, such as VoIP.

2.1. BITAG Interest in this Issue

Internet standards assume all hosts on the global Internet can connect directly to each other, on any specified port number. However, the practical reality is that blocking of Internet port numbers is a technique used by network providers for various reasons, either in the short term while a permanent solution is found or long term when there is no better solution. Some of these reasons relate to network or user security while others relate to business practices. Port blocking has the potential to cause unintended and unanticipated problems for the operation of applications. Its use also has the potential to be anti-competitive, discriminatory, otherwise motivated by non-technical factors, or construed as

---

[1] In this report, whether a port block is considered "inbound" or "outbound" will be in relation to the user. Please note the terms "inbound" or "outbound" are also used in this report to indicate the direction of Internet data traffic, among other things, and when used in such a manner may be in relation to the ISP, user, or application — depending on the context.

[2] Throughout this report, the term "user" may be used somewhat interchangeably with the terms "consumer", or "customer". Please note that "customer" also refers specifically to the individual or entity that is in a contractual customer agreement with an Internet service provider (ISP), while "user" can refer to both customers and non-customers alike.

such.

Concern has been raised that port blocking is an area of confusion for users, and a cause of friction for application developers, as there does not seem to be uniformity as to:

- Why ports are blocked,
- Which ports are blocked,
- Where ports are blocked,
- Opt-out options,
- Disclosure of port blocking policies,
- How such policies may affect application providers and consumers alike

BITAG aims to address some of these concerns by documenting how port blocking works, the rationales behind it, its implications for different segments of the Internet ecosystem, and suggested best practices for entities that implement port blocking.

ISPs may take a different approach to port blocking depending on whom they serve. Enterprise-focused ISPs, for instance, usually do not implement port blocking as enterprise customers generally have greater security expertise and rarely create additional risk for the ISP's network. Consumer-focused networks generally implement port blocking more often, as their customers are typically less technically sophisticated and unaware of security threats and vulnerabilities. As a result, this report will focus on consumer networks.

2.2. What are Ports?

In the architecture of the Internet, communication between two systems is identified by five fields: (1) the source IP address, (2) the destination IP address, (3) the transport protocol in use, (4) the source port, and (5) the destination port used by the transport protocol [RFC793]. The pair of IP addresses representing two systems identifies all of the communication sessions between them, whereas the port number pair identifies an individual communication session.

Transport protocols, most often Transmission Control Protocol (TCP) or User Datagram Protocol (UDP), include in their header fields two numbers in the range from 0 to 65535: the "destination port" and the "source port" [RFC6335]. When an application on one device wants to communicate with an application on another device, it directs the local operating system to open a communication channel (usually called a "connection") between itself and the remote end point, and specifies the IP address (either IPv4 or IPv6), transport protocol, and port number that the service will use. Applications that can use either a UDP or TCP transport frequently use the same port number for each, but this is neither required nor assumed. For further reference, throughout the remainder of this report TCP and UDP ports will be denoted with the name of the transport protocol followed by a slash and the port number: TCP/520, for example.

By convention, most server applications "listen" on a dedicated port number. For example, in a web communication, the browser on a client system opens a TCP connection to a web server using port 80 as the destination port and a random port number as its source port. The web server, listening for incoming communication on port 80 will invert the port numbers in its response to the client's request. Thus, the web server response to the client system uses port 80 as the source port and the source port learned from the client's initial request as the destination port. In this way, computers can efficiently manage multiple sessions between peers, or between clients and servers. While many applications use standardized destination ports, others choose ports at random when they are establishing communications.

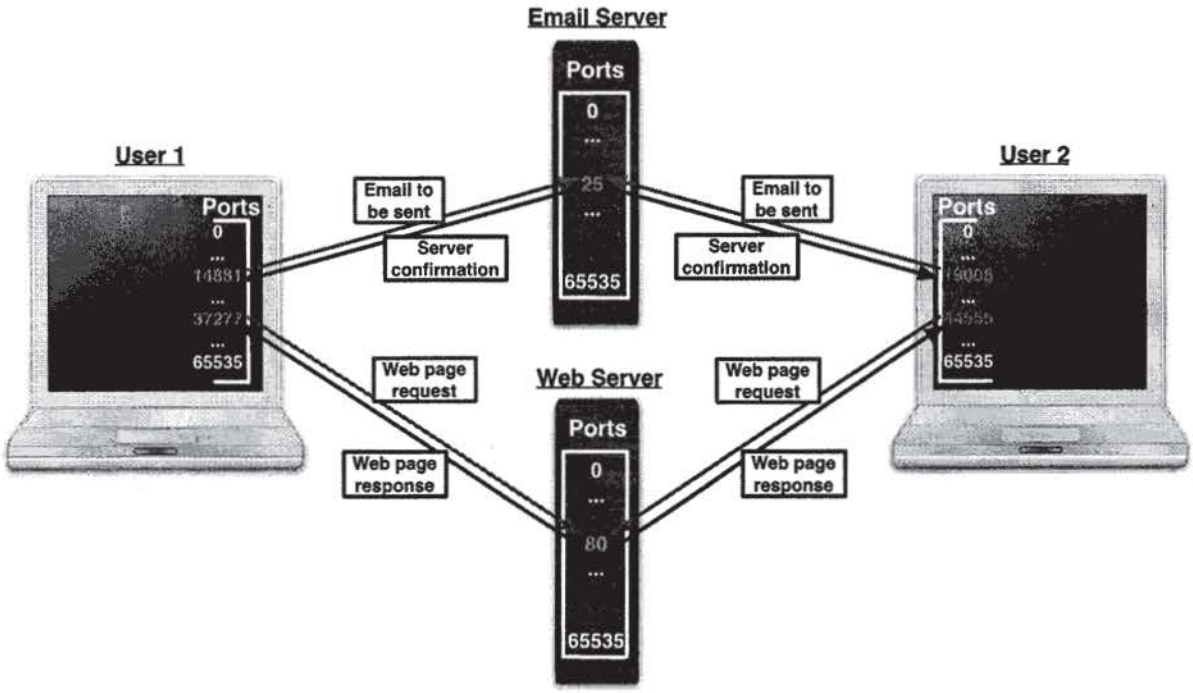

Figure 1: Communications between users and application servers. The users' computers use randomly assigned ports and connect to well-known ports on the servers: port 25 for SMTP email and port 80 for HTTP web traffic.

Port numbers in the range of 0 through 1023 are referred to as Well Known or System Ports [BCP165]. Over time, there has been a need to extend the number of assignable ports. User Ports in the range of 1024 through 49151 are now available for registration of services and protocols through the Internet Assigned Numbers Authority (IANA) [BCP165]. The remaining ports in the range of 49152 through 65535 are referred to as Dynamic Ports. These ports have been set aside for local or dynamic use and cannot be assigned. Client operating systems may use a port from the Dynamic Port range as the source port when originating a request, such as to a web server. The list of assigned port numbers is available from the IANA by accessing the Port Number Registry [Port Number Registry].

The procedure for obtaining a port number or otherwise updating the registry may be found in BCP 165 [BCP165].

It is also possible, and common, to use unassigned port numbers. This happens when an application is in development or is used only in a confined domain, or is "port-agile" in the sense that it is designed to intelligently use any available port number. Port-agile applications may be benign; Skype and other peer-to-peer applications are often port-agile [Skype FAQ]. Traffic created by Distributed Denial of Service (DDoS) attacks and malicious software (known as malware) is also often port-agile or uses a wide range of ports [SANS].

2.3. What is Port Blocking?

As noted in Section 2 above, port blocking is when traffic is identified and blocked on the basis of the combination of transport protocol and port number. Port blocking can be conducted by ISPs, enterprises, or on customer equipment in the home. Because some applications are designed (or were originally designed) to run only over specific ports, a network that blocks those ports prevents those applications from sending traffic unless the applications are redesigned or reconfigured to use different ports.

As an example, the Simple Mail Transfer Protocol (SMTP) was originally designed to use a destination port of TCP/25 [RFC788]. Malware that sends "spam" email frequently does so directly from the infected system to the target, while legitimate email often uses ISP or enterprise email servers as intermediaries. Therefore a common method used to minimize spam is for the access network provider to block traffic from its users that has TCP/25 as its destination port, unless that traffic is directed to one of the ISP's email servers. Users whose email clients are affected by these blocks must reconfigure their clients to use another port.

Port blocking is generally ineffective against port-agile applications or traffic. Applications that use randomized ports or different ports per user or per instance of the application cannot effectively be stopped with port blocking.

Network administrators and home users have a variety of techniques at their disposal for preventing unwanted communications to and from the Internet. For example, traffic coming into an ISP's network might be blocked on the basis of its source IP address (a practice known as Ingress Filtering) to prevent spoofing or to block email sessions that do not traverse the ISP's email servers as a means to prevent spam [RFC2827]. Using a firewall is another technique and provides the ability to block traffic based on different criteria such as source or destination IP address, transport protocol, port numbers, some application-layer criteria, or a combination of these elements. Firewalls come in a variety of types and may be installed on user devices (computers, home routers, etc) or in the network by enterprises or ISPs. Finally, enterprises or WiFi hotspots may prevent all Internet traffic from coming in or out of their networks unless the traffic flows through an HTTP proxy on the network. The most common reason to adopt these techniques is to

prevent network attacks and abuse, although they may be used for other purposes (parental or employee controls or capacity management, for example).

2.4. Network Address Translation (NAT) vs. Port Blocking

Port blocking is distinct from Network Address Translation (NAT), but both can have similar effects on applications. NAT was originally designed to help network operators cope with the scarcity of IPv4 addresses by allowing multiple end devices to share a single public IP address. NAT has been deployed within home networking equipment for years, and is now being deployed with increasing frequency within ISPs' networks (where it is known as Large-Scale NAT or LSN) during the transition from IPv4 to IPv6 [BITAG Large Scale NAT Report].

NAT, by its nature, blocks all unsolicited inbound communication into the network. This is because, with multiple devices sharing the same public IP address, a NAT device does not know which user to send inbound traffic to unless (1) there has been recent outbound traffic using the same address and port, or (2) the NAT device has been pre-configured with a rule on how to map the combination of an external address and port to a corresponding internal address and port. Thus, when a NAT device receives traffic with any source/destination port combination for unknown mappings, then that traffic will be blocked. From the perspective of an application, this effect can be similar to a port block that blocks traffic on the inbound ports the application is designed to use – the application traffic will not reach the user. If the use of LSN continues to become more prevalent, applications may continue to experience these kinds of blockages, without it being obvious to the application provider (or the user) whether the difficulties are caused by port blocking, LSN, or some other functionality in the network or the home.

2.5. IPv6 and Port Blocking

Operationally, network operators have not yet seen widespread security threats or abuse in IPv6 networks, and at this writing there has yet to be significant consumer or enterprise use of IPv6. If needed, port blocking can be implemented in IPv6 as in IPv4, in which case the recommendations of this report apply.

3. Motivations for and Implementation of Port Blocking

Port blocking is a tool commonly used by ISPs, but the use of that tool can vary dramatically from ISP to ISP. Many ISPs use port blocking to protect their customers from security threats, but some have used it to block high bandwidth or competing applications. Some networks block the ports of all of their customers, some allow opt-out and some do not

implement port blocking at all. Understanding the more common applications of port blocking, the differences in customer bases, and how some ISPs implement port blocking will help to illustrate why some ISPs see port blocking as necessary and why to date there has been little uniformity when it comes to port blocking policies and practices.

3.1. Mitigating Known Vulnerabilities or Issues

The ports most commonly blocked on the Internet today are the result of known vulnerabilities in applications running on well-known ports. ISPs deploy port blocking most often as a defense to known security vulnerabilities, easily exploited applications, or as a means to discourage the abuse of legacy protocols when newer standards emerge. The use of port blocking in these instances typically involves a determination by the ISP that the benefit of protecting users or the network from these security risks outweighs any negative impacts upon users. Another common rationale for the use of port blocking is to block traffic unwanted by the ISP's users, e.g., in denial-of-service attacks where a user can be overwhelmed by maliciously generated and unwanted traffic.

Netalyzr is a free web-based measurement tool created and managed by the Networking Group at the International Computer Science Institute that performs network testing and analysis [Netalyzr]. Service reachability, one of the tests included in this tool's suite, attempts to ascertain which ports a service provider blocks by attempting to connect on 25 well-known ports. A report published in 2010 based on 130,000 test sessions showed that four well-known ports are blocked by a significant percentage of broadband service providers. [Netalyzr2010] Those ports are as follows:

Service	Port
SMTP	TCP/25
RPC	TCP/135
NetBIOS	TCP/139
SMB	TCP/445

The rationale for blocking each of the ports listed above is described in turn below, along with the rationale for other commonly blocked ports such as: TCP/161 and UDP/161 for Simple Network Management Protocol (SNMP); other network management ports; and finally TCP/80 for Hypertext Transfer Protocol (HTTP).

3.1.1. Simple Mail Transfer Protocol (SMTP) - TCP/25

One of the best-known uses of port blocking by ISPs is for TCP/25. The Simple Mail Transfer Protocol, SMTP, was originally designed to send electronic mail from one system to another using TCP/25 in an untrusted manner using the Internet Protocol. SMTP is used to send messages from a mail client to a mail server, as well as between mail servers. Mail clients use a separate protocol to retrieve messages from mail

servers. Most users today do not send their email via TCP/25 since more secure mechanisms have evolved.

However, since the original design of SMTP an increasing amount of spam email has been transmitted using TCP/25, often providing false information about the email address of the email client. In addition, malware is often transmitted using TCP/25. Computer viruses designed specifically to send spam in this manner, as illustrated in Figure 2, represent a significant threat to the functionality of the Internet and to ISP operations as well. Service providers are further incentivized to either block or at least monitor for nefarious activity on TCP/25 as a means to prevent their respective customer IP address spaces from being placed on email blacklists. Blacklisting ISP address space prevents customers in that address range from sending mail to a large percentage of email destinations, as many ISPs prohibit the receipt of email from blacklisted source IP addresses.

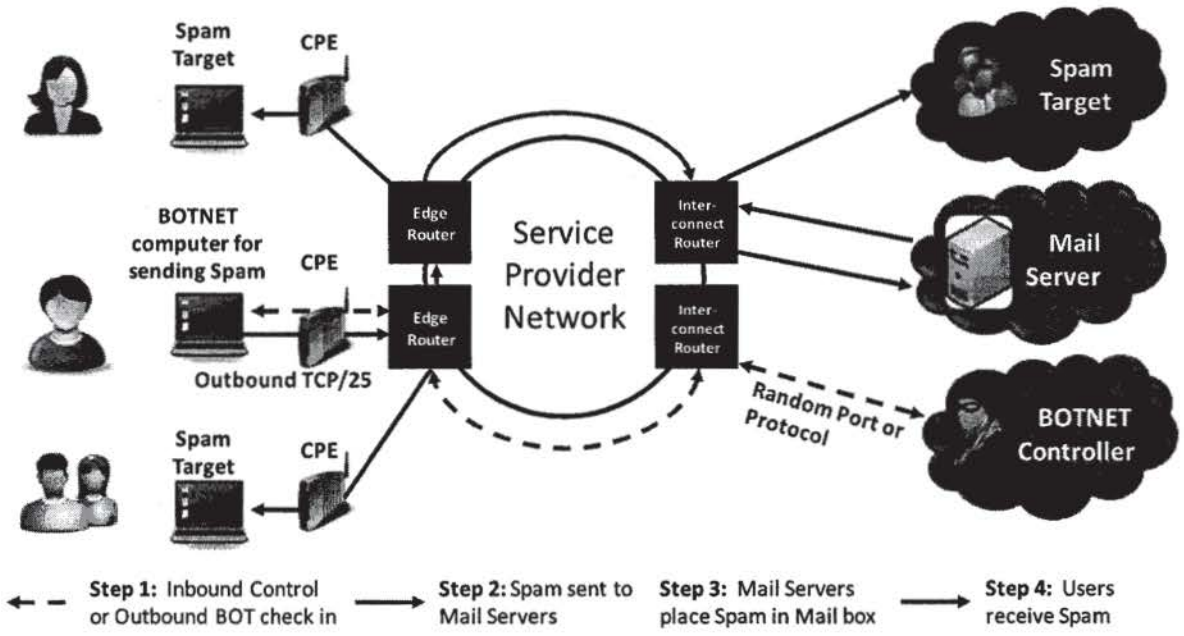

Figure 2: Spam across unblocked TCP/25

To combat these security issues, Internet standards for SMTP have evolved. The standards now support communication using alternative ports, including TCP/587, and support using authentication to ensure that the email client identifies itself correctly and is an authorized sender [RFC6409]. The Internet Engineering Task Force (IETF) further recommends that communication from email clients to mail servers transition from TCP/25 to authenticated TCP/587, and that mail servers should similarly authenticate all email, even if received on a different port [RFC5068]. However, Internet standards continue to rely on TCP/25 for forwarding of email between mail servers [RFC6409].

An industry trade group has gone further by recommending that ISPs block outbound TCP/25 from all machines on the ISP's network other than the ISP's own mail servers and block inbound TCP/25 traffic, which can thereby reduce the transmission of spam from infected computers through mail servers outside the ISP's network [M3AAWG Port 25 Recommendation]. It is important to note that the blocking of SMTP on TCP/25 does not prohibit the customer from sending email. Customers are generally instructed how to configure their mail clients to use the viable alternative ports, such as TCP/587 for email submission.

However, the IETF notes that blocking of outbound TCP/25 can be problematic for some users and that there are alternative established practices for controlling abuse of port 25, including the use of proxies and/or rate limits, and thus offers no recommendation concerning the blocking of TCP/25 [RFC5068]. In addition, blocking TCP/25 not only blocks communication between email clients and servers on TCP/25, but may also prohibit the ISP's users from running their own mail servers.

Most Internet service providers implement outbound TCP/25 blocking adjacent to the customer's point of connectivity to the network as illustrated in Figure 3, in order to prevent sending of spam SMTP mail from customer computers, while some service providers only block outbound TCP/25 for those customers suspected of sending spam. Some service providers may also implement inbound TCP/25 blocking to prohibit SMTP traffic directed toward the ISP's users from sources outside the ISP's network. Some service providers that block TCP/25 will remove this block for individual users upon request (by the customer), while others will not.

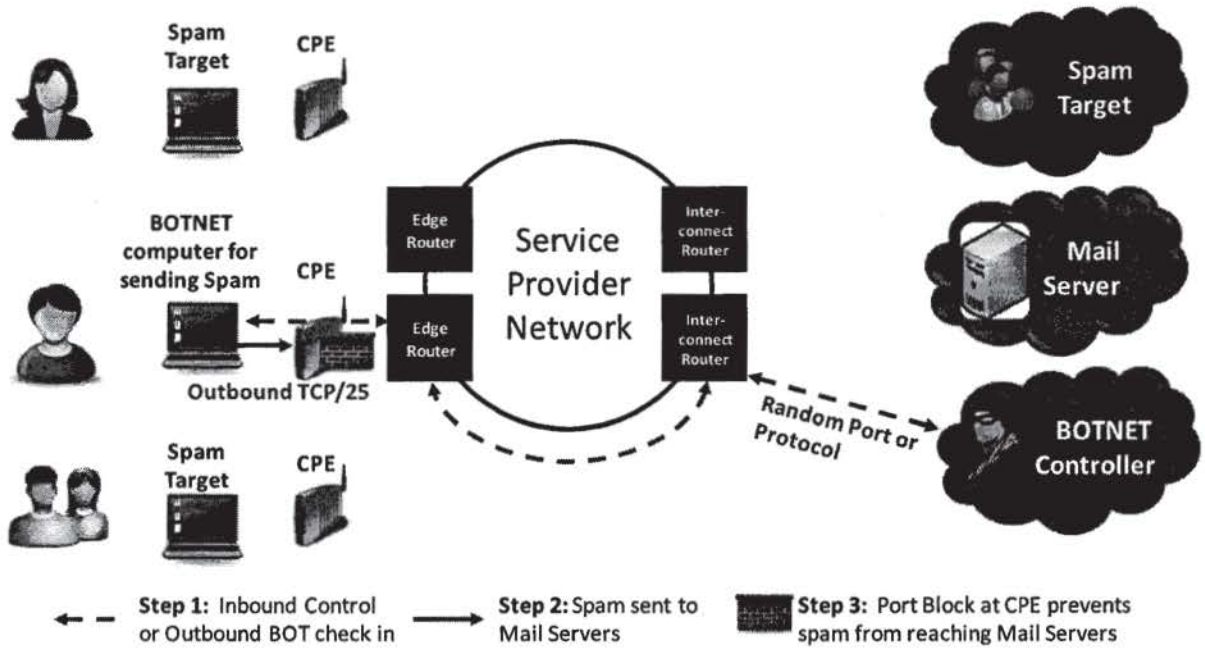

Figure 3: Spam blocked at CPE on TCP/25.
The computer infection is not corrected but rather the negative impact on others is mitigated.

3.1.2. Microsoft RPC – TCP/135 and UDP/135

Another well-known port that many ISPs block is TCP/135, associated with Microsoft's Remote Procedure Call (MS RPC). MS RPC utilizes TCP/135 and UDP/135 for communication between clients and servers and between clients and other clients. Vulnerabilities in Microsoft's RPC code were exploited by a number of large-scale Internet viruses including the Blaster and Reatle worms. The mechanism was also exploited as a means of delivering spam using popup messages. The impact of the Blaster worm in 2003 was large enough to negatively impact a number of ISP networks. As a result many ISPs implemented TCP/135 and UDP/135 blocking as a defensive measure to prevent their customers from propagating the virus.

3.1.3. NetBIOS and SMB – TCP/UDP139 and TCP/445

Two additional well-known ports that many ISPs block are TCP/139 and TCP/445. NetBIOS and Server Message Block (SMB), two services associated with the Microsoft Operating System, use these ports.

The NetBIOS service was originally developed for the IBM Personal Computer and was later adopted by Microsoft's MS-DOS Operating System and subsequent versions of the Microsoft Windows software. NetBIOS was originally a Local Area Network service and later extended to support TCP/IP for Internet operation [RFC1001] [RFC1002]. NetBIOS includes naming, connection-oriented and connectionless services, and promotes communication among trusted network devices including file sharing, printer sharing, etc. NetBIOS utilizes UDP/137 and UDP/138 as well as TCP/139 for communication. TCP/139 applies specifically to the session connectivity service provided in the NetBIOS protocol suite.

The SMB protocol operating on TCP/445 is typically considered in the same category as NetBIOS based on its use in Microsoft Windows software. The SMB protocol is closely associated with NetBIOS for file and printer sharing among a group of computers running the Microsoft Windows OS. Security vulnerabilities have been found with both SMB and NetBIOS that allow remote users to gain control or execute malware on unprotected computers in home networks. Combined with the fact that these protocols are predominately used for communication among devices within the home network and not for sharing of services over the Internet, many operators have chosen to block communication directed to these ports to protect customer computers from malicious actors external to their network. Some service providers block both inbound and outbound ports in these ranges, while other service providers do not block these ports at all. If these ports are blocked, then a user will find that file and printer sharing will be more difficult to accomplish to or from a remote destination. There are alternative remedies to blocking these ports, most notably through operating system patches and through direct control over these ports by each user on each of their computers; however, many users do not apply security patches in a timely fashion and do not properly configure ports on their computers.

3.1.4. Simple Network Management Protocol (SNMP) - TCP/UDP 161/162

Simple Network Management Protocol (SNMP) has been subject to widespread abuse, particularly for amplification DDoS attacks that take advantage of the relative ease of spoofing the source address of UDP packets, and is blocked by some ISPs. Please refer to a recent BITAG paper for more information on SNMP abuse [BITAG SNMP Report].

3.1.5. Network Management Ports

Another category of ports that are blocked by some ISPs support network management traffic that can be considered harmful or deemed inappropriate when originating from a customer, unless explicitly permitted by the ISP. Protocols that fall into this category include routing protocols or network management protocols originating from the customer's equipment and directed upstream toward the ISP's router. Examples of these protocols include Dynamic Host Configuration Protocol (DHCP) and Routing Information Protocol (RIP). DHCP for IPv4 operates on UDP/67 and UDP/68. DHCP for IPv6 operates on UDP/546 and UDP/547. RIP operates on UDP/520. Some ISPs implement port blocks that are strictly designed to prohibit a customer's misconfigured network device from impersonating an ISP's DCHP server. Similarly, ISPs may implement port blocks to prevent a customer's network device from attempting to use the RIP protocol directed at the ISP. In both the DCHP and RIP use cases, the customer's traffic is directed at the ISP's network equipment and not the Internet. The blocking of these ports is unlikely to have a negative impact on users.

3.1.6. Terms of Service Enforcement

Some ports are blocked to enforce an ISP's terms of service. The most common example of this type of port blocking concerns Hypertext Transfer Protocol (HTTP), an application protocol widely used on the World Wide Web to request and to transmit web pages. HTTP was standardized by the IETF and the World Wide Web Consortium, and is most commonly used over TCP/80 [RFC2616]. Blocking inbound TCP/80 prevents a customer from hosting a web page, but does not prevent the customer from surfing the web.

HTTP can use a number of ports, but TCP/80 is the default port and most commonly used. A user's web browser thus usually transmits requests for web pages to the remote server using a destination port of TCP/80, and receives web pages using other ports.

ISPs that block inbound TCP/80 commonly justify the practice as either a security concern or enforcement of their terms of service. The common security justification is that blocking inbound TCP/80 stops malicious traffic, such as the Code Red worm, that attempts to infect a computer through this port. The terms of service justification is that blocking inbound

TCP/80 prevents users from running web servers, which may be prohibited by the ISP's terms of service for consumer-grade Internet access.

However, both justifications have weakened over the past few years. The security concern for TCP/80 can be prevented by proper configurations of the user's firewall and/or computer operating system and thus many ISPs no longer block TCP/80. While many ISPs previously used their respective terms of service to prohibit consumers from running web servers, either to manage upstream capacity or to differentiate consumer Internet access from business Internet access, some ISPs now use different methods to differentiate services. Thus, blocking of inbound TCP/80 has previously been more common than it is today, and most ISPs no longer block TCP/80 [Netalyzr].

3.2. Alternatives to Port Blocking

As the above sections have demonstrated, there are a variety of different circumstances and security threats that may provide motivations for ISPs to institute port blocking. As such, the alternative approaches available for resolving any particular problem will depend on the problem itself. For example, many of the vulnerabilities discussed above could also be resolved by applying software updates, patching operating systems, installing consumer firewalls, or upgrading home equipment. However, these alternative solutions cannot typically be implemented rapidly at a large scale, and in many cases are outside the control of the ISP. Whether consumers or ISPs pursue these alternative solutions, and whether the alternatives effectively reduce the risks that could otherwise be mitigated by port blocking, depends on the ease-of-use and costs of the alternatives.

3.3. Persistent Versus Temporary Blocking

Port blocking can be implemented temporarily on a short-term basis or persistently for an extended duration. Temporary blocking is usually done for security reasons, with the block typically remaining in place until the security threat is eliminated or substantially lessened. The vulnerability can be eliminated or reduced over time through software updates to the vulnerable devices, changes in technology or elimination of depreciated equipment. The time scale seen for temporary blocks ranges from hours to weeks.

In some cases it may prove impractical or impossible to remove these temporary port blocks. The examples of port blocking detailed above in Section 3.1 are temporary tactical blocks that evolved into persistent blocks, with durations now measured in years. These persistent blocks address known vulnerabilities that, for various reasons, either cannot or will not be corrected in the near future. Some of these vulnerabilities can only be addressed through fundamental architectural redesign of the Internet.

3.4. Considerations for Different Types of Networks and Network Architectures

As noted in Section 2.1, ISPs may take a different approach to port blocking depending on the type of customers they serve, or the type of network they run. It is important to understand that ISPs implement port blocking to different extents on different types of networks.

Enterprise-focused ISPs generally implement little to no port blocking due to the technical sophistication of their enterprise customers, while consumer-focused networks generally implement port blocking more often, as their customers are typically less technically sophisticated and unaware of security threats and vulnerabilities.

For a number of reasons, cellular operators are more sensitive than wireline operators to attacks and abuse that create large traffic volumes. As a result, some wireless carriers may be more aggressive than wireline carriers in their use of security mitigation tactics, including port blocking, because: (1) cellular data networks generally have less available bandwidth than wireline broadband networks, due to the limitations of radio; (2) wireless devices have limited battery power to expend on the additional processing required to defend against traffic created through attacks or abuse; and (3) wireless customers are more often subject to usage-based billing plans, where they are charged for the amount of data used.

3.5. Where do Port Blocks Occur?

Port blocking can be implemented at many different places in the network path. In a residential network, the most common places are typically located at the:

(1) Service Provider's Network Interconnection Links between ISPs
(2) Service Provider's Customer Facing Network Links
(3) Customer Premises Equipment (CPE)

With communications occurring in two directions (to and from the device), blocking can also be directional. Port blocking policies can, and often do, distinguish between inbound and outbound traffic. Since the session is identified by a pair of addresses and a pair of port numbers, and traffic from a client to a server always uses the identifying port number as its destination port, it is straightforward to prevent sessions in one direction while permitting them in the other.

3.5.1. Network Interconnection Port Blocks

The first common location port blocking can be implemented in an ISP's network is at the network interconnection links to other ISPs. Blocking inbound traffic at this location removes the ability of sources outside the ISP's network to send traffic on these ports to the ISP's users. Blocking outbound traffic on specific ports at this location removes the ability

of the ISP's users to send traffic on these ports to destinations outside the ISP's network. However, neither inbound nor outbound port blocking at this location removes the ability of the ISP's users to send traffic on these ports to other users of the same ISP.

From an ease of management perspective, this is the best location in the ISP's network for implementation as it is the quickest to deploy because it requires the fewest number of interfaces to provision and manage. Blocking at the network interconnect links effectively protects against external threats because it impacts all of the external traffic.

As shown in Figure 4 however, the implementation of port blocking at these locations does not protect the ISP's customers from one another. In addition, these links are also typically large capacity links, and while today's routers are capable of implementing port blocking without a performance impact, historically that has not always been the case.

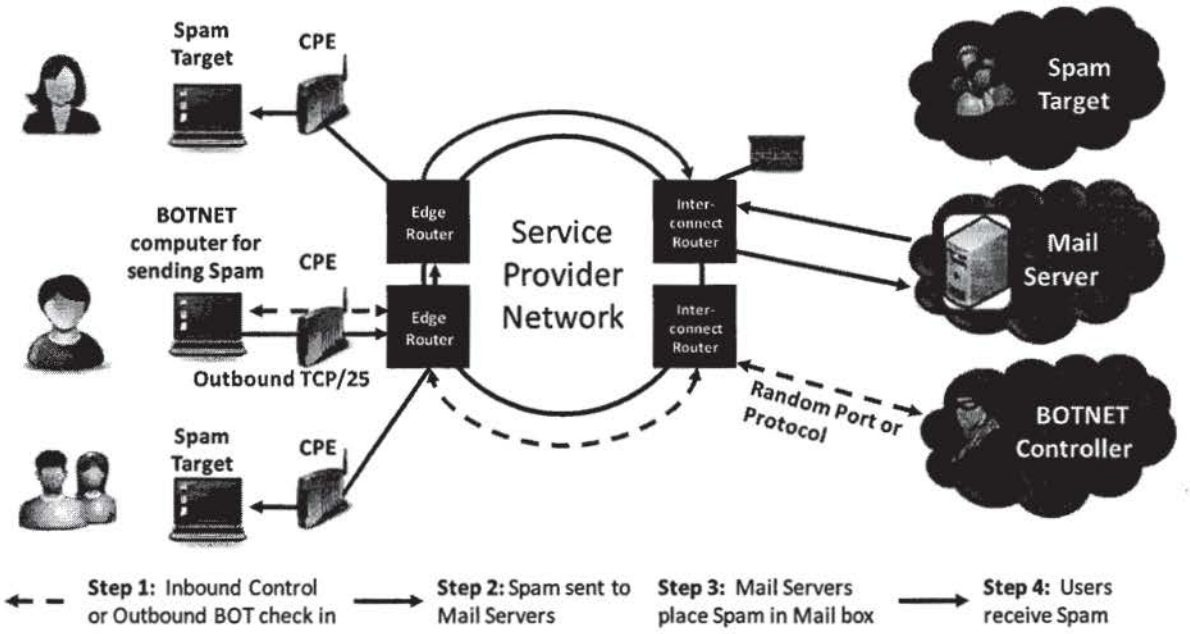

Figure 4: Spam blocked at Interconnect Router TCP/25.
The computer infection is not corrected and is allowed to impact users on the Service Provider network. Port blocking on the Interconnect Router only helps targets that are not service providers.

3.5.2. Customer Facing Network Connection Port Blocks

The second common location to implement port blocking in an ISP's network is at the customer aggregation routers on the customer facing links. The effect of blocking traffic at this location varies depending on the type of network. In some networks inbound port blocking will prevent any traffic on that port from reaching the customer, while in other networks it will allow traffic from only other customers in that local area. Likewise in some networks an outbound block will prevent all outbound traffic on that port and in others will still allow outbound traffic on that port to other customers in that area. In contrast to

the first location described above, port blocking at this second location requires blocking on substantially more interfaces.

Management of the port blocking policy is more complex and time consuming at this location due to the increased number of interfaces, but it has historically had less (or no) impact on the performance of the network as these interfaces are lower capacity. If port blocking is intended to address security concerns, then the primary reason many ISPs have chosen this location to implement port blocking is that it provides more protection to customers than the Interconnection link location, and provides good-to-adequate protection against malicious customers that would harm or burden the network. This location allows the ISP to both protect the customer from threats on the Internet and protect the Internet from threats from customers, as seen in Figure 5.

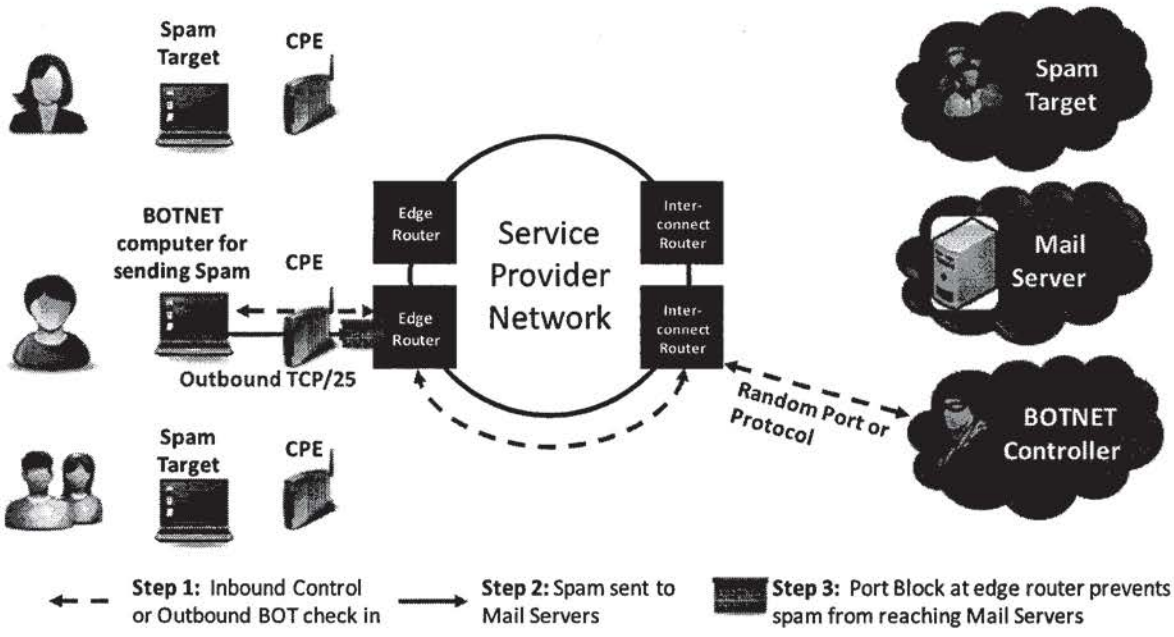

Figure 5: Spam blocked at Edge Link outbound TCP/25.
The computer infection is not corrected but the spam is prevented from reaching the mail server.

3.5.3. Customer Premises Equipment Port Blocks

The third location where port blocking is commonly implemented in an ISP's network is within customer premises equipment (CPE), e.g., cable or DSL modems and/or home routers or gateways. If inbound port blocking is implemented in equipment at the customer's premises, this removes the ability of sources outside the customer's premises to send traffic to the customer on these ports. Blocking outbound traffic on specific ports at this location removes the ability of an ISP's customer to send traffic on these ports to destinations outside the customer's premises. None of the ISP implemented blocks affect traffic within the home.

Port blocking on customer premises equipment distributes the processing load required to implement port blocking. This location is also the most granular of the three common locations, in that it allows an ISP to apply port blocking rules on a per customer basis. However, at least one trade-off is the high cost of administering port blocking rules on what could potentially be millions of devices rather than a few thousand interfaces.

The viability of this third location can vary depending on the type of equipment located at the customer premises, but especially on the ownership or "control" of the equipment – as some equipment is provided by the ISP and some by the customers themselves, with different levels of control or ownership by each group. Below are two example scenarios:

Scenario 1 – CPE is managed by the customer

In this scenario, the ISP does not provide a home router or customer premises equipment to the customer, or the ISP provides the device but does not manage the security policy on the device. Many times, the customer will purchase a home router that typically implements a certain level of port blocking by default. However, since the ISP does not manage the equipment, the ISP will not have the ability to implement port blocking at this location. From the perspective of the ISP, this scenario presents a high risk to both the customer and the network. From the perspective of the customer this offers the highest amount of freedom, but also implies that the customer adopts the associated risk.

Scenario 2 – The ISP provides a device that is capable of providing port blocking and is solely managed by the ISP

In this scenario the ISP maintains the ability to control the port blocking policy (protecting the customer and preventing malicious customer traffic) while still retaining the flexibility to modify the rules in some instances on a case-by-case basis if requested by the customer. From the ISPs perspective this affords all the protection of port blocking in the network. From the perspective of the customer, this scenario may reduce the risk to the customer, but may also reduce the freedom accorded to the customer if the ISP does not allow opt-out from port blocking.

In general, if a customer manages the CPE device in his or her home, the ISP has no ability to provision port blocking rules in the device and may therefore resort to implementing port blocking in the network if no reasonable alternatives are available. In cases where the ISP manages the CPE and has the ability to apply port blocking rules in the device it may allow the ISP more flexibility in implementing port blocking policies and opt-out.

3.6. Opt-Out Options

As illustrated in the previous sections, there are several technical variables to consider for when and where a port block is applied, and whether or not an opt-out solution is offered to users. Beyond the technical variables, additional consideration is given to: (1) financial concerns, (2) operational factors (overhead of managing the potential multitude of individual policies), (3) Internet reputation (allowing cyber attacks to originate from the service network can impact all of the ISP's users), and (4) legal concerns (risk of applying the wrong policies to the wrong people).

The impact of an ISP's port blocking policy on the user will vary according to which ports are blocked. In some cases, e.g., blocking of network management ports, the port blocking policy is unlikely to have a negative impact on the user even if opt-out is not allowed. In other cases, where the port blocking policy may negatively impact some users, allowing for an opt-out policy could help to minimize any negative effects.

Due to the enterprise and commercial customers' high level of technical sophistication, these customers often are trusted by their ISP to connect to its network without the need for port blocking rules set by the ISP. In a similar manner, there can be a certain percentage of residential customers with the same level of technical sophistication – who may warrant a similar approach.

The decision of whether and how to allow a user to opt-out of some or all of the port blocking rules within an ISP's port blocking policy may depend on the ISP's rationale for the block, where and how the port block is implemented, and most importantly the ISP's network design. The capability and the cost of implementing an opt-out option vary greatly from ISP to ISP. Some ISPs will find opt-out technically impossible, some ISPs will find opt-out possible but costly, and some may find opt-out relatively easy. While some of the factors in this decision are highlighted in this report (such as where the blocks are implemented), every ISP will face its own specific complexities (such as IT design). For these reasons some ISPs allow for opt-out, some require the user to move to a business service (which are designed to not use port blocking), and others do not allow opt-out at all.

4. Implications and Concerns Relating to Port Blocking

The implications and concerns related to port blocking may depend upon where a stakeholder "sits" in the Internet ecosystem. An ISP may see port blocking as an indispensible tool while an application developer may see port blocking as a challenge. Users may have different perceptions or concerns based upon their level of technical expertise. This section will offer some of these differing perspectives, as well as touch on some of the security considerations that go along with port blocking.

4.1. Concerns of Internet Ecosystem Stakeholders

4.1.1. Internet Service Providers

The Internet was built around the premise of an open and shared environment. Many early Internet protocols were designed with limited or no security measures built into their basic communications. Today, applications that leverage these protocols have inherited their minimal security characteristics.

Many ISPs implement port blocking to address some or all of the long-term problems discussed in Section 3.1. If left unblocked, these are threats that can cause an increase in spam or can compromise users' information. From the perspective of most consumer ISPs, the implementation of these port blocks can dramatically reduce support costs (less customer calls, less spam complaints, etc.) and result in minimal or no inconvenience to most users.

Most ISPs will also use port blocking as a means to mitigate a short-term, or what is hoped to be a short-term, threat. These threats generally fall into the denial of service (DoS) category. Some good examples of this are SNMP to prevent SNMP-amplified Distributed Denial of Service (DDoS) attacks, and some worms like Blaster. While port blocking is not a silver bullet, and does nothing to prevent port-agile attacks, it can be a viable short-term mitigation step for some attacks until a long-term solution is found.

Port blocking is used by a few ISPs to enforce terms of service. Some ISPs offer different levels of service such as consumer and business and can differentiate those services by allowing one to host servers and another that cannot. The number of ISPs that use port blocking in this fashion has declined in recent years, and there are only a few ISPs left implementing this practice.

Historically, there have been instances of ISPs using port blocking to block bandwidth-intensive or competing applications. While these are the implementations that draw the most attention, ire, and press, they have been very rare.

Consumer ISPs serve a variety of customer types ranging from a majority of customers who are less technically sophisticated to a minority of 'power users'. Unfortunately, there are a number of bad actors that connect to the Internet and a large number of customers that do

not adequately protect themselves. ISPs must walk a fine line of trying to create a service that will protect the average user while not hindering the power user. In general, ISPs believe port blocking is a critical tool in order to create the secure environment the average user wants, and that port blocking for security reasons rarely causes problems for application developers or power users.

4.1.2. Applications and Application Providers

Ideally, application providers would be able to design their applications under the assumption that the entire port number space is available to them. However, port blocking, together with NAT, firewalls, and other technologies, generally serves to reduce the number of available ports.

Port blocking can complicate application design and development and create uncertainty about whether applications will function properly when they are deployed. Although port blocking may be intended to block only unwanted traffic, it may also inadvertently block wanted traffic by mistake. This inadvertent blocking of wanted traffic may lead application developers to move their applications to ports that are not blocked. However, the availability of any particular port – outside of the "well-known" ports – can be somewhat uncertain, as different network operators can independently choose to block or unblock individual ports. Thus, applications designed to function across multiple IP networks must take into account the potential for port unavailability or unreliability (although applications may need to do so anyway, perhaps more so because of the prevalence of proxies, firewalls, NAT, and LSN rather than ISP-based port blocking).

Some application providers may be concerned about ISPs intentionally blocking their applications for anti-competitive purposes. For example, in 2005 the ISP Madison River was found to be blocking ports associated with independent voice over IP services that were in competition with the ISP's own voice telephony services [Madison River]. Port blocking has also been used for the purpose of limiting traffic from applications associated with high traffic volumes, such as peer-to-peer file-sharing applications [Toward Quantifying Network Neutrality]. Such concerns have caused some application developers to adopt some of the mitigation measures discussed below, such as designing their applications to be port-agile or using ports unlikely to be blocked.

Whether and how an application provider chooses to mitigate the effects of port blocking will depend on a number of factors, including the size of the impact on the application's user base, the expected duration of potential blocks, and the rationale behind the blocks. A short-term block may not require mitigation, whereas blocks that affect many users and are expected to remain in place over the long term may trigger more extensive responses. If blocking is conducted for non-technical business reasons, application developers may choose to contest those port blocks in business negotiations, regulatory forums, or in public rather than developing mitigations.